Chapter 8

The Control Subsystem

The illusion of freedom.
The feeling of being under control.
Who is controlling who?

8.1 Introduction and Motivation

In considering a natural or human made system that operates as expected, it is often difficult to differentiate between process and controller. Is the teacher in control of the class, or is the class in control of the teacher? Even simpler, when flushing the toilet there is the clear understanding that the action of starting the process, be it by a proximity sensor, or pressing a button, or pulling a chain achieves the desired effect: the toilet is cleaned, and ready for its next use. Nobody really cares about the internal workings of the toilet system. From this perspective, it is difficult to realize that a feedback loop is at work, let alone that certain references such as flushing time, and fill level have been set.

In some other cases, despite the fact that control and process are well-integrated it remains relatively easy to identify what is in control and what is controlled.

A steam engine in a locomotive is such an example. There is the combustion chamber where the fuel is burned, producing steam in the boiler, which through its pressure and expansion drives the crank and slider mechanism that turns the wheels. There is also a Watt's governor[1] as shown in Fig. 8.1. This device opens or closes the steam supply according to the actual wheel speed, too fast and the steam supply is throttled, too slow and the steam supply is opened up. It ensures that the locomotive runs at a nearly constant speed, as prescribed by the driver.

In this case it is very easy to distinguish two subsystems: the energy transformation chain from the chemical energy in the fuel to the kinetic energy of the train, and the governor

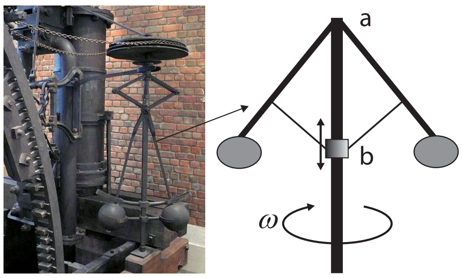

Fig. 8.1. Steam locomotive and Watt's governor

[1] The centrifugal force on the balls turning around *a* due to the speed of the train raises the point *b* and reduces the steam supply. James Watt, 1736–1819, Scottish engineer whose role in the industrial revolution cannot be overestimated. The unit of power is named after him.

regulating the speed. Steam engines would not run safely without a governor, neither would wind mills for that matter (it is from the latter that Watt adapted the idea to suit his engines).

The structure of this example is repeated in natural as well as engineered systems.

Control is now so pervasive that it is hard to consider any device that does not contain some form of control. Hard disk drives and compact disk players have track seeking and servo control. Cameras have image stabilization. Washing machines have sequencing, and often rule-based controllers that minimize energy and water usage as dish washers do. Microphones and hearing aids and recorders have automated gain control to avoid saturation and improve the fidelity of capturing sound. Cars are full of control devices, some have brake-by-wire, traction control, cruise control, climate control, engine control, lights and screen wipers that respond to the driving conditions, adaptive suspension and so on. Similarly aircrafts, trains, ships … Nowadays driven by a need for greater efficiency or more sustainability there is a push for pervasive networking, creating smart infrastructure (water, electricity and gas distribution systems as well as transport networks) and developing the largest networked control systems spanning entire continents.

So far we have considered the simple interconnection of systems as *open loop* or *closed loop*. This chapter goes somewhat beyond this, and is particularly concerned with the control subsystem, what it consists of, and what we may expect from it, and how it is interconnected in the overall system.

Historical Note

In its early incarnations, feedback was an art, an inherent part of making things work. The earliest examples of feedback are probably found in the first mechanized ancient water clocks (clypsydra) found in ancient Rome and Greece, as early as 300 BC. The water clocks, which essentially keep time through a constant flow of water, were made more precise due to a pressure regulator in the water supply. Another well-known example of feedback art, again related to time keeping, is the design and indeed evolution of the escape mechanism, essentially a speed regulator that brings the swinging motion of the pendulum over to the motion of the hands of a mechanical clock. All of these were implemented by trial and error, without much theory or prediction (which comes with analytic design) of how they would function.

The trial and error approach failed at the dawn of the industrial revolution. Indeed, the steady state oscillations or *limit cycles* that appeared in some applications of steam engine governors lead to the first systematic study of stability, with contributions by Maxwell[a], Routh[b] and Hurwitz[c]. Similar governors were in use well before the steam engine, to regulate the speed of wind wills all over England and the European mainland. They functioned well, but the interaction with the much faster working steam engine created dynamic instabilities not encountered before. It took about from the middle of the 19th century to the middle of the 20th century to transform the art of feedback into an engineering discipline. The introduction of digital computing enabled automation to proliferate, and now control engineering is a normal part of engineering curricula in most engineering disciplines all over the world.

[a] James Clerk Maxwell, 1831–1879, best known for the Maxwell equations describing the electromagnetic field. One of the greatest minds of all time, according to Einstein.

[b] Edward Routh, 1831–1907, English mathematician, best known for his work on mechanics, and an analysis of stability of linear systems.

[c] Adolf Hurwitz, 1859–1919, German mathematician which become famous for his work on algebraic curves.

The role of the control unit is intimately related to how it is interconnected with the object under control. Its position in the cascade or feedback loop conditions how the control acts in the overall system, and this in turn affects how it is designed. The actual nature of the control unit can be mechanical, electrical, electronic or digital, depending on the most appropriate technology and the application. The different elementary ways in which the control subsystem can be linked with the system are shown in Fig. 8.2.

Control System Interconnection

a *Open loop, control before the plant*. The control unit is to generate the appropriate command inputs for the plant.

b *Open loop, control after the plant*. The control unit evaluates the output of the plant and decides about future inputs for use at a later stage. Typical in batch processes and so-called quality control, not for real time control.

c *Feedback loop, feedback control*. The control unit generates the inputs for the plant, based on real time information about the response of the plant, and information external to the plant.

More complex structures can be conceived, but they all use the above as elementary building blocks.

These different options will be further explored in Chap. 10, where we discuss design aspects of the control system. The purpose of this chapter is to analyze the structure of the control subsystem itself. What is it for? What are the options? What can we expect from it? We already know the main concepts to be discussed in this chapter. They are the information flow, the control objectives and the control constraints, feed forward versus feedback, mixed control, integrating process and control.

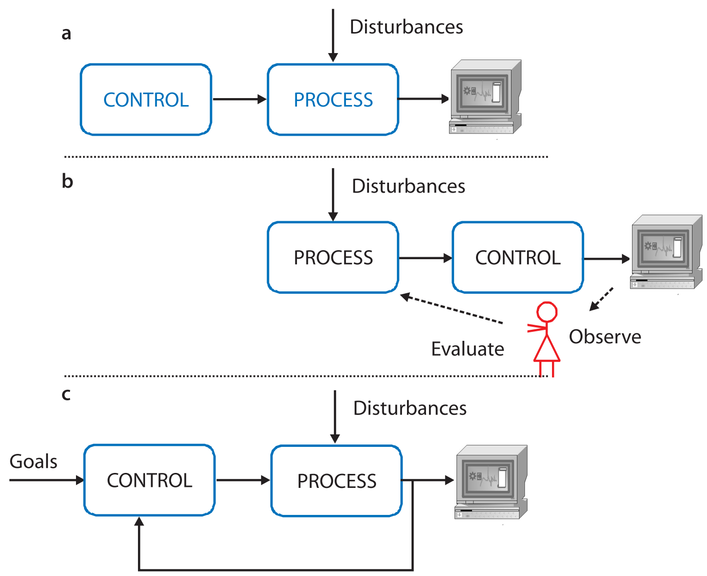

Fig. 8.2. Process and Control subsystem interaction

8.2 Information Flow

At a basic level, it is clear that monitoring and control both depend on information extracted from the process (perhaps through sensors) and that control units deliver information into the system for further action by the process (perhaps through actuators). Causality determines the direction associated with the information flow.

The basic structures of information flow are a cascade and a loop. So far we have only seen some simply interconnected systems, either a single cascade or a single loop, and occasionally a cascade containing a loop like in the two degrees of freedom controller structure. Also it has been always easy to determine the causality between subsystems, and hence the information flow was always easily recognized.

However dealing with a natural system of some complexity, or even an engineered system, it is not always easy to tell how the information flows. Nevertheless, causality can always be inferred from an understanding of the dynamics. Moreover, when the topology of interconnections becomes cluttered, with many interacting loops and cascades mixed, even simply keeping track of what controls what becomes a major undertaking. The rule here is really divide and conquer. First identify the information flow for each interconnection: what is the cause, what is the effect (even this may be hard if the subsystems at the interconnection cannot be isolated from the rest of the system). Then using the resulting block diagram, subsystems in cascade can be studied separately, or in a single block. Any loops will need to be considered as a single unit. This way we can build a hierarchy of levels of systems, until we arrive at the external signals, the free inputs, and the measured outputs. At the end of this process, we can zoom in and out analyzing the system at any one particular level of its subsystems, and thus building our understanding of the overall behavior.

In a process environment, the Fig. 8.3 schema is quite typical, at least at the highest level. There is the main stream of information from the manipulated variables to the measurements gathered by the Data Acquisition System. The process generates signals

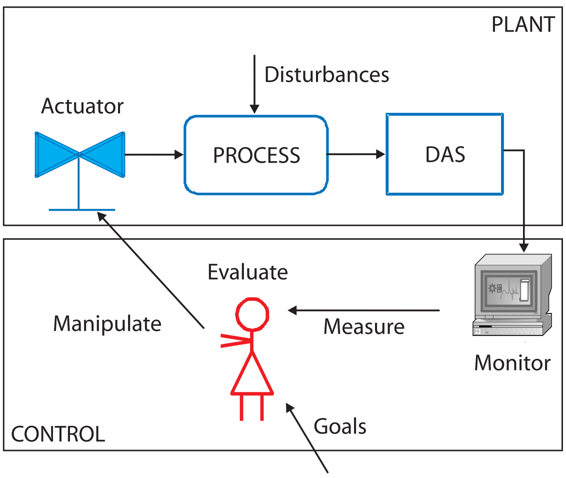

Fig. 8.3. Information flow

which provide the information about its behavior, as influenced by external disturbances. This information is then digested, and presented to the ultimate controller, the human operator who decides where to go next with the system.

This picture at this level of generality even applies to how a country's economy is run. An important collection of operators sets the price of money (the interest rate), which in turn influences the economy, the vital data of which, like trade balances and debt levels and unemployment conditions allow the operators to make further decisions.

In case of a distributed control system, the flow of information is much richer. There are many subsystems intertwined, all with different dynamics, acting over different time scales. Understanding the information flows is essential to come to terms with how control acts in a system.

8.3 Control Goals

No matter at which level we look into a system, the main purpose of a control subsystem, the controller, is to generate signals that eventually (perhaps through human intervention) drive the process it controls, such that the latter behaves in an expected manner. There is always at least one physical interface between the control subsystem and the process, and in a feedback loop there are two, one for sensing and one for actuation.

In some cases the interface is just a direct physical connection. For instance in the Watt's governor (see Fig. 8.1) the displacement of the flyball(s) is mechanically linked to the displacement of the valve position that throttles the steam flow.

In many other instances, the output of the controller exists in a totally different physical domain from the domain in which the input signal to the plant lives, or the sensed variable lives. In such instances a transducer translating signals between physical domains is essential to realize the interconnection. For instance, the controller computes a digital representation of the next input, which must be converted into the appropriate process variable, like the position of a gate, a voltage across a motor winding, or a force or the speed of a motor or the intensity of a light beam and so on. In engineered systems the transduction typically involves the use of electronics, and electrical actuation through motors that then link into the domain of interest. Similarly in the human body, the electrochemical representation of motion in the neural system must be transformed into a muscle force before motion actually occurs. This happens through a chemical process at the neuromuscular junction: a motor neuron's stimulated synapse releases neurotransmitters that bind to a receptor at the muscle fiber which then reacts by contraction.

Based on the objectives and the knowledge about the process (its model, and other relevant information like what are the typical disturbances), as well as the direct measurements, the controller computes the subsequent appropriate actions for the system to react to.

Some common control objectives are

- *Regulation or disturbance rejection* as illustrated by the governor in the steam engine. The goal is to keep the shaft speed at a preset level in spite of disturbances such as a change in the available steam pressure due to changing conditions in the burner. In the irrigation system (see Sect. 3.3) the water level in the supply channel had to be regulated, so that farmers have the right amount of potential energy available to

supply water across their fields. The disturbances are due to the variable water demand on the channel.

- *Tracking*, to follow a reference, like what we saw in the radio telescope example (See Sect. 3.4). The goal is to track a star whose position relative to the observatory is changing over time. External disturbances, like the wind load, must be rejected. Tracking often comes with disturbance rejection requirements as well.
- *Sequencing predetermined procedures* as is required in the start-up or shut-down phase of a process. This happens typically before a process reaches its steady-state, where regulation or tracking becomes the main aim. For instance, before the speed regulator is able to control the speed, the steam must reach a sufficient temperature and pressure. Starting up a steam engine from a cold condition, requires a sequence of events: fill the boiler, check the water level, start the burner, open the air valve, open the fuel valve, initiate ignition, check combustion, reach operational conditions. A similar procedure is to be followed to shut the machine down. Such sequencing, and emergency procedures exist in most processes. In the irrigation system, channels have to be filled at start up. Rain events lead to an emergency shut down. The radio telescope must be stowed under heavy wind conditions to avoid damage. In other applications, like in batch reactors or the simple washing machine or the car wash, sequencing is actually the main control task (see applications described in Chap. 3).
- *Adaptation* such as maintaining overall system behavior despite significant changes, it typically requires large scale adjustment of the control system. Since the dynamics under control change, it is appropriate to also adjust the control subsystem, to ensure that its suggested actions remain appropriate. In the antennae system (Sect. 3.4) the mechanical resonant frequency changes with the pointing angle of the antennae. This affects the allowable bandwidth of the control input, and the controller is adaptively tuned to cater to this.
- *Optimization*, rather than regulating a variable to a specified level, sometimes there is a need to optimize a variable: like in attaining maximal efficiency, or highest power output. This can be seen as regulating the derivative of the variable to zero. Specific control algorithms for such tasks have been developed, like extremum seeking methods. These are closely related to adaptive control and/or learning control methods.
- *Fault detection and process reconfiguration*, as from the monitored process behavior it is feasible to identify alarm conditions, and take action through the control system. In this manner unsafe operating conditions can be avoided either automatically as may be demanded in an emergency, or in advisory capacity to suggest possible actions to an operator. Sometimes, an alarm may require the need to reconfigure how control acts, which is particularly the case when actuators fail. In combination with physical redundancy of actuators, a control system can reconfigure how it implements future control based on the fault detection and/or alarm condition.
- *Supervision* is typical in situations where there are many levels of control, as in the exploitation of a large utility network. The network changes as operating conditions vary, or the network structure itself changes (switches), or components fail. An evaluation of the new condition will typically lead to an assessment of which control objectives remain valid or not, which resources are available for control, and based on this assessment the lower level control subsystems are redirected on how to act. At this level large scale simulation and scenario assessment take center stage in control design.

- *Coordination* is a task normally executed at the highest level of control being typical when there is a clear hierarchy of control subsystems. Coordination ensures that the various subsystems are properly working together. Local control subsystems are directed and provided with references or set-points. When dealing with large scale complex processes, each subprocess will have its own goals, that are aligned with the overall goal. The coordination ensures that the local goals are attained in such a way that the overall process objectives are met. It also assists in start-up and shut-down as well as emergency procedures.
- *Learning* can be achieved from signals gathered during the operation of the controlled system. New information about its behavior can be discovered, and learning or adaptive control techniques can capture this information for later use, like to optimize the controlled response. Learning in itself, in particular when used in feedback leads to very complex system behavior, which is still not well understood, as even in the simplest examples of such control systems chaotic behavior cannot be excluded.

All these different objectives result in very distinct control approaches and techniques: from logic and event-based, to discrete-time controllers, to sophisticated intelligent control systems where ideas from artificial intelligence are key. All these control systems emulate to some extent control as we experience it from our own behavior, just as Norbert Wiener argues in Wiener (1961).

These different flavors of control come with their own analysis and design tools. In modern applications multiple tools are used cooperatively to arrive at an acceptable design solution. Thus far the theory and practice of control engineering has not evolved to the point where there is a standard software "*Control Engineer version 4.13b*" that will guide someone to an appropriate control solution from the problem's inception.

Example

Let us re-consider the steam generator (Fig. 7.8) with the human in-the-loop, as in Fig. 8.4.

The main goal is to produce steam to meet the demand, of course all the while keeping the plant within its proper operational envelope. That means, an appropriate water level, appropriate fuel burning conditions, safe pressure levels, and so on. One aspect of the control unit is to ensure that the water level and water temperature are properly *regulated*.

The operator monitors the overall working of the plant, via a display panel, as in Fig. 8.4. The operator can track the evolution of the system and tune some of its settings. He or she can program the plant to produce a different steam product (changing the pressure, temperature, flow rate) depending on the requirements, either scheduled or as to meet demand. The operator can introduce a *sequential procedure*, to start the boiler up or shut it down, or go into a stand-by mode.

In a more complex operation, the control unit could *supervise* many of the operating conditions and decide automatically about changes in the plant or in the control. For example, based on economic pricing of the fuel, the burners could switch between gas or oil; or the control objective could be switched from regulating the boiler in stand-by to tracking under variable steam load conditions, even *adapting* the subsystem controller that takes care of the water level control, as the dynamics change due to the increased

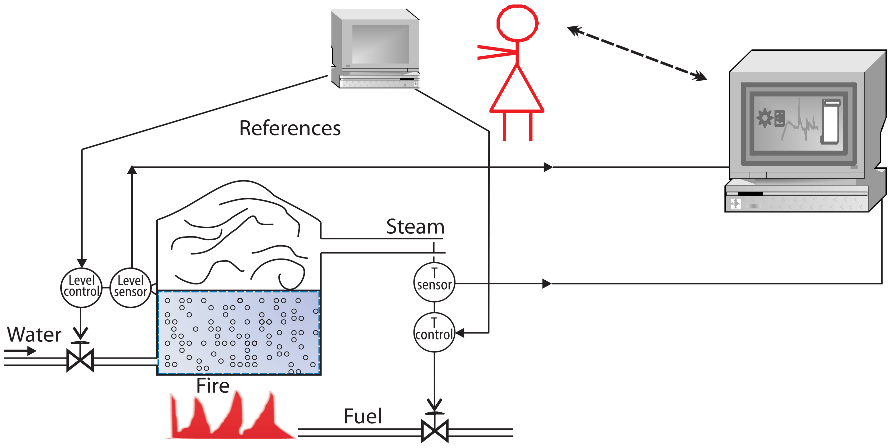

Fig. 8.4. Boiler water level and temperature control with an operator

steam flow. Controllers, subservient to the main controller, can aim at maximal efficiency in the burner, within the limits imposed by environmentally acceptable flue gases.

The control system must be able to *react to faults*. For instance, an emergency shutdown of the burners when the water level is critically low, so as not to physically damage the boiler plant.

Similar situations can be foreseen for all systems described in Chap. 3.

8.4 Open-Loop

As already mentioned, open-loop control refers to control structures where the information flows in a unique sense, without closing a loop. The user determines the parameters of the controller to be applied, the controller generates and applies the control signals to the process, the process dynamically evolves, subject to possible external disturbances and, finally, the user evaluates the results.

In this case, we still have a closed loop, in that the operator closes the loop, but there is no machine-only information loop.

Several open-loop situations can be distinguished.

Sequencing

Consider a washing machine, or a car wash system, or a simple CD player:

- The equipment is well understood. The components, their operation and all the options are clear.
- The process is well-known in advance. Everything has been done before, all events, normal and abnormal are catalogued.
- Disturbances are not really expected, but the process can be engineered to be foolproof (engineers do try).
- The performance requirements are not very strict.

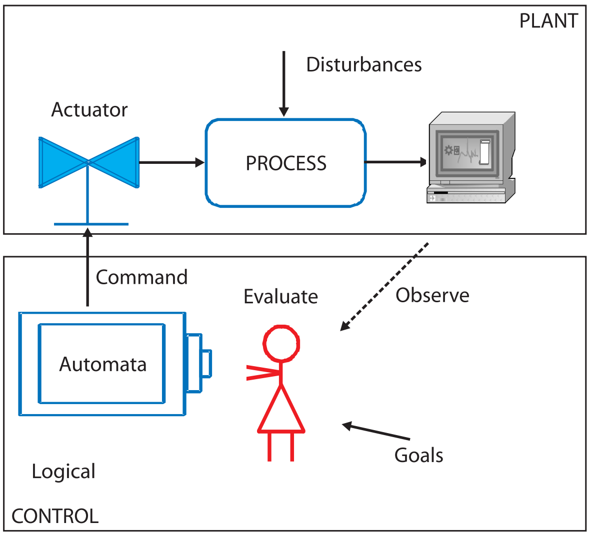

Fig. 8.5. Sequential control

There are many other applications of the same ilk, like starting a car engine, air conditioning, heating a home, operating a video camera. In all these cases, automata provide the appropriate language to deal with the analysis and design of such systems, as shown in Fig. 8.5. Automata are programmed, and then execute a sequence of tasks as events unfold. They may react to specific events (in which case they do use some feedback, and are only open loop in between event times), or simply be time-based in which case they are truly open-loop and entirely oblivious to the unfolding process.

Open-Loop Control

Open-loop means that the manipulated variables are controlled without taking into account the actual evolution of the plant, so there is no information loop.

This is what happens in most master/slave systems. The master generates the control actions (probably in a sophisticated way) and the slave is just applying these signals to the slave system. A key duplicator is a mechanical master-slave system: the "sensor" sweeps over the original key profile and the cutter is tracking this signal on the raw key. Probably we all had the experience that not all the duplicated keys work well at first! If the error is because the cut is too superficial, a second pass can polish the result, but if the error is due to an excessive cut, the created item is useless.

More critical are the master/slave applications familiar from tele-operation. If the application is a surgery, the patient is at risk. If the application is for sampling rocks on Mars a great deal of investment, effort and time is at stake.

The general structure of a computer-based open-loop control system is similar to that of the sequencer. The actions to be applied are often computed in advance, stored in the controller and applied at a later time. The main drawback is indeed the lack of (immediate) feedback. If the actual system's behavior is not as expected, it deviates from the model

used to compute the control actions, or if there are some disturbances not taken into account in this computation, the actual system's response will not match expectations.

On the other hand, if the operational conditions are precisely those used to generate the control signals, the system will perform without any error. The controller does not need the existence of an error to generate the control, as it is the case in feedback control. The overall system is easier to implement. Most robots used in manufacturing operate in open loop.

Feedforward

The control signals can be generated on-line in real-time based on information gathered from sources other than the system itself. *Feedforward* control applies to an open-loop control structure where the control action is computed based on measurements of disturbances or references, as shown in Fig. 8.6.

Let us consider again the boiler system. In this case, we focus our attention to the level regulation in the water tank. As soon as a user opens a valve to get more steam, we do not need to wait until the water level decreases. We know that steam consumption requires more water and hence in response to increased steam demand the controller opens the water inlet valve (and may also increase the burner's heat output).

Sometimes like in multivariable systems, the disturbances on one controlled system are actually created by the actions from another control subsystem that controls another variable that interacts with the first control subsystem. Again in the boiler example, if we try to control simultaneously the water level and the temperature we know that a desired increment in the temperature will need an increment in the heat output of the burner. This modifies the equilibrium water/steam, the pressure and also the

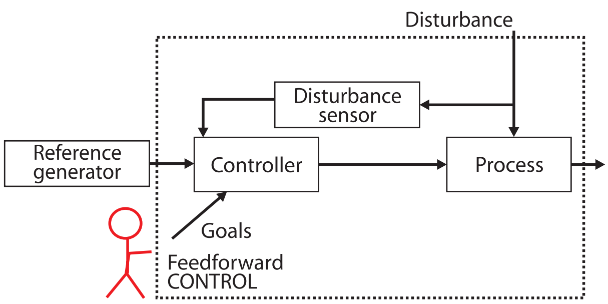

Fig. 8.6. Feedforward control

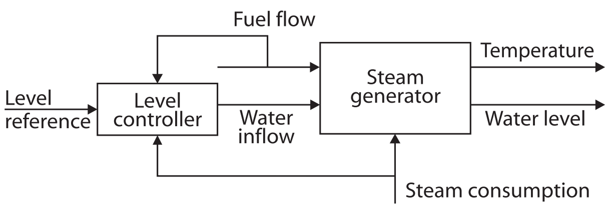

Fig. 8.7. Feedforward control of water level

water level. This interaction must be taken into account to reduce undesirable actions from the water inlet valve. These ideas are illustrated in Fig. 8.7.

Similarly, when tracking a reference signal we are able to know this reference signal in advance so we may exploit this preview knowledge to ensure that the system response is well-aligned with the reference at all times. We use this strategy when driving a car, we look ahead not in the rear view mirror to steer the car. Driving a car by feedback only would equate to driving using information only from the rear view mirror and other instruments in the car.

Even so, feedforward is rarely implemented without a feedback loop of some sort, alarms must be monitored, and goal attainment is to be verified.

Quality Control

A special case of an open-loop control system goes by the name of *quality control*. In this methodology, quite common in manufacturing industries, measurements, evaluation and corrective actions are taken after the completion of the entire process.

For instance, consider the quality control of the ceramic tiles in Chap. 3. The procedure involves

- measuring the tile characteristics (size, defaults, appearance, …);
- classifying the tiles against predefined standards;
- statistical analysis of the data, leading to a proposal of future actions, perhaps changing process variables or using new set points in the control subsystem;
- reporting the findings to management.

In general, these activities are performed off-line, and often the reaction time is actually longer than the production time of a single batch of product.

In order to gain more from the quality control observations, some of the activities could be fast tracked, to approach real-time feedback. Often this uses heuristics and/or rule-based actions on partial measurement results. In this way corrective actions can be applied in a shorter time frame. The investment is worth it if there is a sufficient economic return because of improved quality or throughput in the production line.

Open-Loop Control

In open-loop control information flows in one direction, without feedback. It works well if:

- The process is well-characterized.
- Disturbances are not important or they can be adequately measured and counteracted.
- Performance requirements are not very stringent.

It is certainly not appropriate if:

- There is uncertainty about the process behavior.
- There are unknown and/or unmeasurable disturbances.

The dynamic behavior of the whole system is readily deduced from the dynamics of the subsystems.

8.5 Closed-Loop

Closed-loop control uses response information to determine the input. Feedback is advised when some or all of the following conditions apply:

- Disturbances are unavoidable.
- System response information is available. This is the *sine qua non* for feedback.
- Performance requirements are demanding.
- System dynamics are poorly known, or subject to large uncertainty, or variable over time.

In feedback, Fig. 8.8, control signals are based on all available information, process knowledge, which includes a model for the uncertainty about the process knowledge, disturbances information (measurements, and/or models), objectives to be satisfied (reference tracking, regulation and so on) and last but not least the actual system response measurements.

What infrastructure is required for feedback? Obviously, there is always a cost associated with feedback. Instrumentation is essential, both for measurement and actuation. The answer to the above question really depends on the overall requirements. The main issue is simply this, the feedback must be able to decide if the requirements are met or not, and have enough action capacity to enforce the requirements.

Sometimes simply the consideration of the objective is all that is required for feedback. For instance, the room temperature regulation by a thermostat is simple: heating on when too cold, heating off when too hot. The thermostat receives a temperature reading that can be rather crude: temperature below or above or inside the reference band is all that is required.

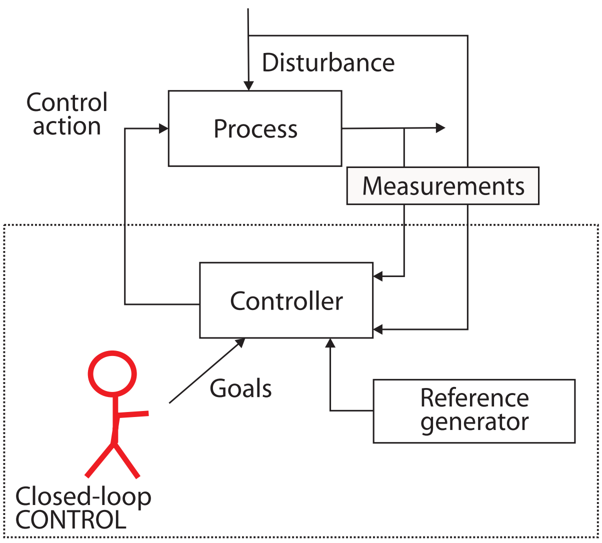

Fig. 8.8. Closed-loop control

In general though, more and more precise information leads to better feedback. There is of course a rule of diminishing returns. Importantly, all measurement information that feedback could possibly need is captured in the notion of a state for the system. A system state is defined as sufficient information at a given time to be able to predict the system evolution given the present and the future inputs. If we can find a *minimal* (smallest dimension) state, that is enough for feedback. There is simply no more information that can be gathered about the system. Hence *state information* and *state feedback* is the most we ever can do.

It sounds great, but state feedback requires in general a lot of instrumentation and a lot of communication capacity to transmit the data. More often than not, state feedback based on full state measurements is prohibitively expensive or cannot be implemented because the state cannot be measured (e.g. try to measure the temperature *inside* a bauxite smelter).

In such situations, and knowing how powerful the state information actually is, the closed-loop control may consist of two separate systems: first a system that takes all available information (measurement, model etc…) and produces from this a (best) guess or estimate for the state (or any desirable internal variable) and next a controller subsystem that uses the estimate of the state to produce the next input value. The first subsystem, from measurement and model to state, is also called a *virtual sensor*, nowadays incorporating many features related to the quality of the measured signal itself. The celebrated *Kalman filter* is such a virtual sensor; this is discussed in Sect. 9.2.

An alternative, using partial state information, is discussed in the next section.

Closed-Loop Control

The basic control loop requires a sensor, an actuator and a controller. Their interaction with the plant to be controlled defines the global behavior of the controlled system.

Closed loop control is essential when:

- the process has to operate in an unstable regime;
- there is large uncertainty about the process behavior;
- there are unknown disturbances.

Feedback must be tolerant of some level of errors, otherwise it can never act. Feedback must be provided with measurements that either directly indicate, or allow the inference of how well the control objective is met.

Feedback may be inappropriate when:

- the required instrumentation is too expensive or does not exist;
- there is no plant uncertainty, nor disturbance;
- feedback design cannot be verified.

8.6 Other Control Structures

Besides feedforward and feedback control many other control structures are used in practice, but they can be seen as a combination, or perhaps a repeated combination of these two basic concepts.

Using a few simple examples, we review some of these other control structures.

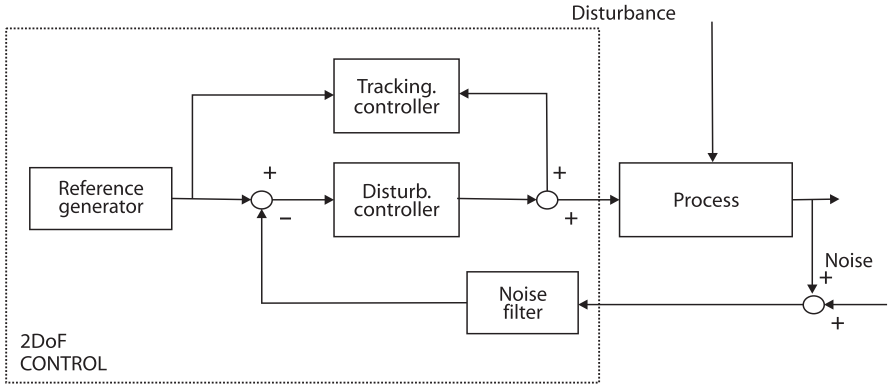

Fig. 8.9. 2DoF (two degrees of freedom) control

8.6.1 2DoF Control

The stability properties of a controlled system are determined by the control loop(s). When stability is of concern, and there are also disturbances acting on the system, and/or there are tracking requirements, a single loop control strategy does not provide enough design freedom to achieve all objectives. Because signals can be shaped by systems, the tracking performance clearly depends not only on the loop but also on any system in cascade. This reasoning leads to the very common control structure with two degrees of freedom, as represented in Fig. 8.9.

The antenna system described in Chap. 3 is a good example of such a controlled system. The antenna should reject the effect of external forces, such as the wind. Also, the antenna must point to an object in the sky that moves (star, satellite). The servosystem is a two-degree-of-freedom controller that ensures sufficient damping of wind effects and ensures at the same time accurate tracking.

8.6.2 Cascade Control

Full state information may be unavailable, or difficult to use at once. In cascade control, successive control loops are used, each using a single measurement and a single actuated variable. The output of the primary controller is an input to the secondary and so on. A judicious choice of how to pair variables and the ordering of the multiple loops can lead to a very efficient control implementation without the need for a full state feedback control.

Again the servosystem for the antennae is a good example. In the end the antenna needs to point to the right position. The position error can be used to set the reference for the motor, but the motor speed and its torque are also important variables (that describe in part how the system works, and they are part of the state of the system). The motor speed can be measured, and so can the current (which is related to the torque) and these can then be used with advantage in controlling the overall behavior. See Fig. 7.3. In this way a multiloop or cascade control system, as depicted in Fig. 8.10, can be implemented.

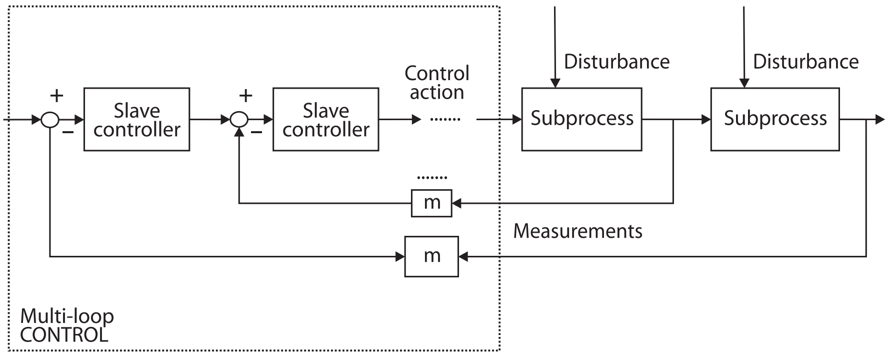

Fig. 8.10. Cascade control

The inner loop will regulate the motor current, effectively ensuring that the motor behaves as a torque source. It eliminates load disturbances. Then, the speed feedback will provide damping and stability. In the outer loop, the position controller (the master controller) will take care of the final position of the shaft. A nice aspect of why this works so well, comes from the fact that the loops have different reaction times. The inner loop is fast, about ten times as fast as the damping loop, which is itself a factor of ten faster than the outer position loop. In bandwidth terms, the outer loop has the smallest bandwidth, the damping loop's bandwidth is roughly ten times larger, and the inner loop another factor of ten larger again.

Thus, the main concept behind the cascade control is: *feedback any information in the plant as soon as it is available*.

8.6.3 Selective Control

Let us consider a waste water treatment plant. All pollution must be removed, and the cleaned effluent should meet all required safety and environmental regulations: its temperature, pH, biological contamination and turbidity must be regulated.

Under normal operational conditions, the pollution controllers determine the maximal effluent flow rate, so that the effluent meets expectations. However, in an emergency, like flood conditions, the controller will temporarily "forget" about pollution control and will determine the outflow to avoid a massive plant failure that may be caused by overflowing storage tanks.

In combustion control dealing with both fuel and air flow, safety will always override normal combustion control to generate heat for the boiler.

Also, to avoid explosive mixtures in a burner, typically air should be in excess. Therefore if there is an increment in power demand, the air flow is increased in advance of the fuel flow. In the opposite direction, if the purpose is to reduce power, the fuel flow is reduced first, followed by the air flow. Thus depending on the required plant evolution either air or fuel tracks the other one.

An example of selective control is also illustrated in Chap. 3, talking about the purpose of exercise (see Fig. 3.26) in deciding the insulin regime for a diabetic.

8.6.4 Inverse Response Systems

As already mentioned, one of the basic features of feedback control is reaction: first an error is detected, then follows the (re)action.

In many systems the error response at the onset is in the same direction it will settle into the future, as long as there are no corrective actions applied. Unfortunately, there are systems where this is not the case. The initial response goes in the opposite direction of where the final response will evolve towards.

Examples are systems with a pure time delay. Actually these do not have any response at all in the initial stage.

There are also the so-called non-minimum phase processes[2]. In these circumstances, the initial response will mislead the feedback controller about what to do. In fact, these systems are much more difficult to control and there are inherent limitations in what control can actually do, no matter how hard we try.

A typical example is the water level control in a boiler. Clearly to keep the water level at a given reference point, the inlet water flow is to be increased if this level drops below the reference. Under normal operation, when the outlet steam flow is increased, the boiler pressure reduces. As a consequence, more water goes to boil, creating more air bubbles inside the water mass, and this has the effect of expanding the apparent water volume in the boiler. The water level sensor observes an increase in the water level. This would require a reduction in the flow of the inlet water, and yet the opposite is required. A typical response is shown in Fig. 8.11. An increment on the steam outlet flow at $t = 6$ s leads eventually to a final reduction of the water level of around 4 cm, but in the first two seconds, the water level actually increased by more than 1 cm. A similar problem occurs when increasing the inlet flow, as this relatively cold water enters the boiler this cools the water mass in the boiler, bubbles disappear and the water level goes down.

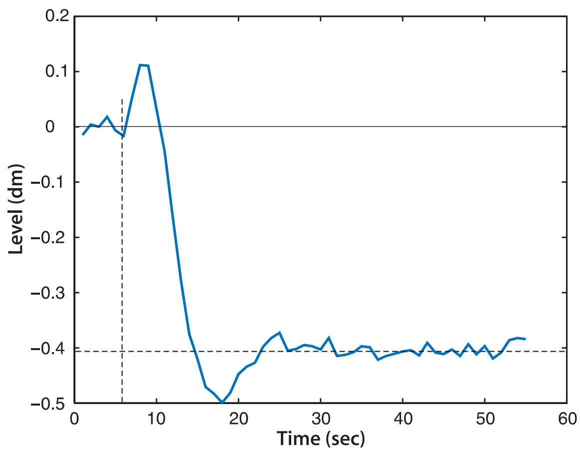

Fig. 8.11. Boiler water level response when the steam outlet flow increases

[2] In the context of linear systems, a non-minimum phase system has a transfer function that has zeros with a positive real part i.e. there are unbounded inputs for which the system has no response at all.

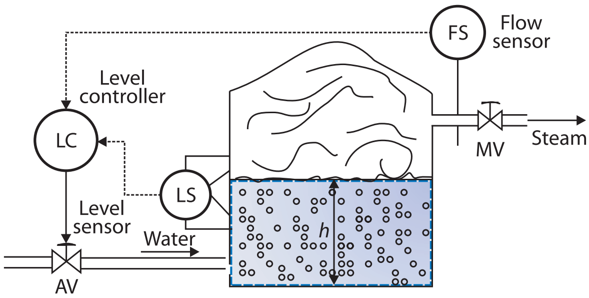

Fig. 8.12. Boiler water level control under steam flow load disturbances

For these kinds of systems feedforward control may help. For instance, in the previous example, if a steam flowmeter is installed, the water level controller can know in advance the forecasted effect of steam flow variations (due to changes in the manual valve MV) and modify the water inflow accordingly (acting on the automatic valve AV) preempting a level error reaction. The measured changes in the steam flow will be introduced as changes in the water flow reference. Nevertheless, as there will be always some difference between the steam and water flows, a final more slowly acting feedback compensation will be required. A typical control lay-out is depicted in Fig. 8.12.

8.7 Distributed and Hierarchical Control

In nature and in engineered systems, many system properties are actually distributed over space, and not characterized by a single or a few measurements. Consider, for instance, the temperature control of the human body, as depicted in Fig. 8.13. First the required body temperature is not the same for all parts of the body. Our heart and brain require a different temperature from our feet and fingers. Moreover many different controlling mechanisms are in operation. In the end though, all these temperature control subsystems are interconnected and must work together to achieve a common purpose. Also there are many different temperatures: skin, muscles, viscera, the brain. There are temperature sensors everywhere as well as end effectors, to react. Some actions are elaborated locally (spinal cord) and some others go through the brain for decision making.

In human made control systems the same rule applies. In a complex system, although the global goal could be the same (optimize the consumption, maximize the benefits) each part of the system has its own control options and local goals. Nevertheless, like in the human body, the communication channels will allow a coordination of controllers to achieve the global goal.

Thus, in a distributed control system we may find simple on/off controllers, automata, simple or sophisticated local controllers and, a lot of information exchange between subsystems to coordinate all the local activities.

Remember the manufacturing of tiles as introduced in Chap. 3. The real objective is to produce as many quality tiles as possible at a low cost with minimal pollution. The system is totally distributed with a number of local controllers for different subprocesses. The command signals for all the subprocesses must be coordinated to achieve the overall objective.

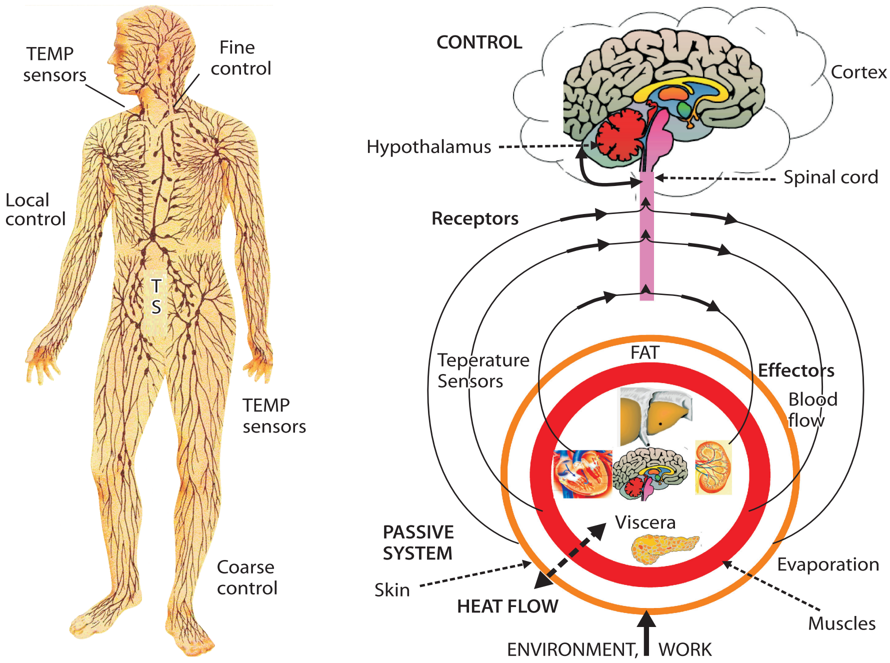

Fig. 8.13. Body temperature control system

Nowadays many distributed control solutions are available. In fact this aspect of control is spurred on by the emerging technology, of sensor/actuator networks. The designer can select components, a communication network technology, appropriate computer control hardware, and appropriate software, and design the system to act in unison. The distributed control task can be conceived as a decentralized system (loosely cooperating local controllers) or as a hierarchy of different levels of control or a mixture. To know what is the actual information topology as well as what information is available for which actions, is in fact an untractable problem in its full generality. Heuristics, experience, and expert design are called upon to come to a robust well-performing and coordinated system.

CCC: Communication, Computation and Control

The new information era is grounded on three pillars:

- Computers to deal with information (hardware and software)
- Communication to make information accessible (channels, emitters and receivers)
- Control to design what to do with information (algorithms)

In this way, networked control, as applied to distributed control problems, is based on the technological advances in communication, computation and control. Sensors and actuators equipped with a (wireless) networking enable feedback. This feedback must be designed to utilize the (always) limited resources of communication and computation to achieve the maximum benefit in system behavior. Without control design, communication and computation leads to a glut of data without information content.

8.8 Integrated Process and Control Design

In the past, it was normally the case that a control subsystem was appended to an existing process to improve its dynamic behavior. As a result, the only option for the controller is to "select" the best process outputs, and inputs from whatever is available. If the process is not designed to accept a particular well-suited input or the variable of real interest is simply not accessible these opportunities are lost to the controller. Many of these constraints are avoidable when the process and controller are conceived together.

Recently a new paradigm has emerged: integrated design of process and control. There is an important interplay between process and control design. Small changes in the conception of the plant can make its control much easier and hence stronger performance may be achieved. In fact, any change in the process design influences the process dynamics and hence how its control is approached.

Conservative process designs are often an obstacle to achieve better results in controlled operation. For instance, an aircraft designed for ease of manoeuvring at high speed may well be unstable at low speed, and hence cannot be piloted. In the aircraft design phase this is then rejected as an unsuitable design. But if the control is designed at the same time, it could be feasible to design a feedback controller that allows the pilot to fly the plane both at high and low speed, because the feedback controller can deal with the instability at low speed. The result is an improved overall performance. A similar example of such integrated design was briefly touched upon in Sect. 7.7. Intuitively manoeuvrability and stability are contradictory concepts in mechanical systems, requiring strong performance on both accounts really demands an integrated control/process design approach.

In exploiting the use of an exothermic chemical reactor, it is often the case that the best outcome is obtained if the reactor is operating around an open-loop unstable equilibrium. An appropriate control system (with a careful safety net, ensuring that there is a way of coping with a failure in the control subsystem) will allow one to exploit this possibility.

Let us conclude our discussion of controller-process design with two more illustrative and classic examples from process control.

8.8.1 Scaling the Process and Its Control

One of the most challenging problems in process control has been the control of the pH of a large volume/flow of say water. This is mainly due to the extreme range of hydrogen concentrations (the scale extends over 14 orders of magnitude) and the extreme sensitivity with which we can keep track of hydrogen concentrations in water. To make matters worse, typically the point for pH regulation corresponds to maximal sensitivity. Extreme sensitivity of the response to the input at the point of regulation is not a good situation to begin control with.

As a result, it is almost impossible to get the right pH in a large tank, it would take forever to control. A practical solution is to scale the process and the control, as shown in Fig. 8.14

In option **a**, the effect of the disturbance (acid flow into the tank) is observed by the pH sensor (probably with some inherent delay) and the pH controller commands changes in a base flow to neutralize the acid. In option **b**, the same is done in two stages.

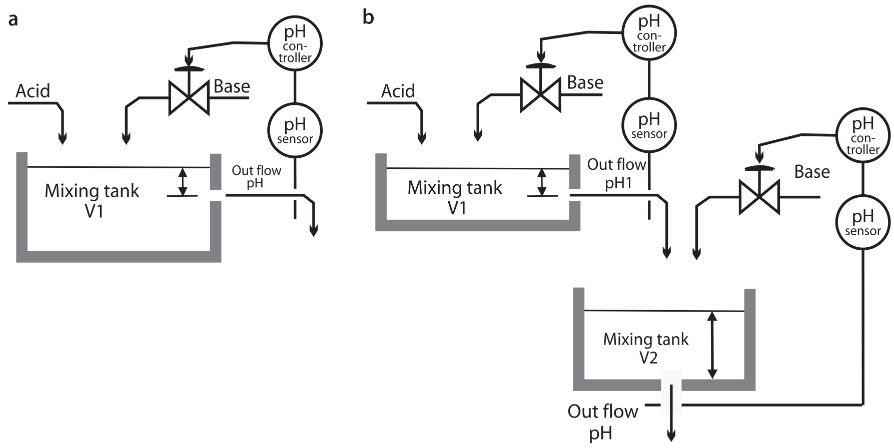

Fig. 8.14. Scaling the process and its control

Dealing with a strong acid (pH < 3) and a strong base (pH > 11), a large control tank requires the patience of Job to settle, as well as an impossibly large actuator with an impossible fine resolution. Not feasible.

Using option **b**, using smaller vessels, in the first tank (a faster and) coarse regulation is achieved using a coarse grained large actuator. In the second tank, the final goal is reached using a smaller actuator with a finer resolution.

Clearly, the process needs to be designed with control in mind. It is not possible to achieve good regulation with a single tank process.

8.8.2 Process Redesign

Similar to the example of the manoeuvrability of the aircraft, we illustrate the design of a typical distillation process unit with two subsystems. The distillation process is split into two parts. If they are conceived independently, the heating/cooling systems are designed and controlled separately, as shown in the upper part of Fig. 8.15a. On the other hand, the system's efficiency is higher and the control easier if both systems are designed jointly and the control of the cooling subsystems is integrated, as shown in part b of the same figure.

Intertwining Process and Control

When designing a new system, it pays to consider the co-design of control and process. The benefits are large when:

- All performance expectations are expressed in advance;
- Control and process are subsystems that are co-designed to meet the performance expectations;
- Feedback is explicitly catered to;
- Maximal design freedom in both control and process subsystems is used to, explore performance limits.

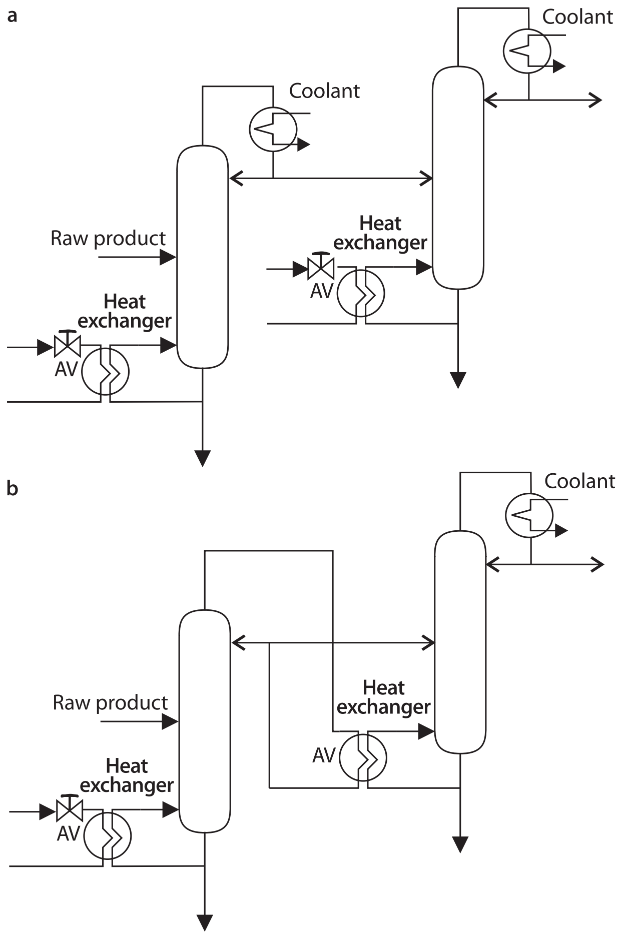

Fig. 8.15. Integrated process and control design

8.9 Comments and Further Reading

In this chapter the basic concepts related to control have been summarized. Feedback plays an important role, but clearly it is only part of the story.

It is perhaps timely to remind ourselves that in feedback we are in the first instance designing with models, not the real thing. Therefore it is important to realize that models are in some sense convenient lies, and we cannot get carried away with the performance of our models and simulations. As always, if it looks too good to be true it probably is.

There are a large number of texts dealing with various design and optimization problems in control. Books dealing in-depth with the process-control interaction typically are about transducers. These tend to be specialized for particular domains, electrical or sound or piezo-electric transducers for example, or deal with a particular application domain. Indeed domain knowledge is extremely important in this context.

Kawaguchi and Ueyama (1989) deals with control in the steel industry. A substantial piece of work just about how the body regulates temperature is Hoydas and Ring (1982). Control and a good understanding of the process must go hand-in-hand to achieve the best possible outcomes, e.g. McMillan and Cameron (2005) deals exclusively with pH control. This is also observed by considering the many different commercial control service providers, they tend to cluster their services in a particular area or domain of expertise. The same can be said for how control design is often taught in disparate manners in mechanical, electrical, chemical or biomedical engineering curricula. Perhaps it should not be so.

A substantial text dealing with control system design in general is Goodwin et al. (2001). In the process control industry Shinskey (1996) is very popular. A book advocating the integrated design approach is Erickson (1999). Multi-loop design is treated in Albertos and Sala (2004) and Skogestad and Postlethwaite (1996). The list is truly immense and the interested reader will easily find additional pointers.